

Application No.:	10/821,372	§	Examiner:	Eland, Shawn
Filed:	April 9, 2004	§	Group/Art Unit:	2188
Inventor(s):		§	Atty. Dkt. No:	5181-94901
Landin, et al.		§	Confirm No.	1210
		§		
		§		
		§		
Title:	MULTI-NODE	§		
	COMPUTER SYSTEM	§		
	WITH PROXY	§		
	TRANSACTION TO	§		
	READ DATA FROM A	§		
	NON-OWNING MEMORY	§		
	DEVICE	§		

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.

I. REAL PARTY IN INTEREST

The subject application is owned by Sun Microsystems Inc. An assignment of the present application to the owner is recorded at Reel 015207, Frame 0686.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences known to Appellant.

III. STATUS OF CLAIMS

Claims 1-32 are pending. Claims 1-32 are rejected under 35 U.S.C. § 103(a). It is these rejections that are being appealed. A copy of claims 1-32 is included in the Claims Appendix attached hereto.

IV. STATUS OF AMENDMEMNTS

No unentered amendment to the claims has been filed after final rejection. The Appendix hereto reflects the current state of the rejected claims.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a system (*See e.g.*, FIG. 20, #100) including a node (*See e.g.*, FIG. 20, #140A) and an additional node (*See e.g.*, FIG. 20, #140B). The node includes an active device (*See e.g.*, FIG. 20, #142A, 146A), a system memory (*See e.g.*, FIG. 20, #144A), and an interface (*See e.g.*, FIG. 20, #148A) interconnected by an address network (*See e.g.*, FIG. 20, #150A) and a data network (*See e.g.*, FIG. 20, #152A) that is separate from the address network (*See e.g.*, specification page 11, lines 18-20). The additional node is coupled to send a coherency message, which requests an access right to a coherency unit, to the interface in the node via an inter-node network (*See e.g.*, FIG. 20, #154; specification page 63, lines 7-15). In response to the coherency message, the interface is configured to send a first type of address packet on the address network if the global access state of the coherency unit in the node is the modified state, and a second type of address packet if the global access state of the coherency unit in the node is not the modified state (*See e.g.*, specification page 92, line 24 - page 93, line 5). In response to the second type of packet, the system memory is configured to send a data packet corresponding to the coherency unit on the data network, regardless of whether the system memory currently has an ownership responsibility for the coherency unit (*See e.g.*, specification page 101, lines 11-15).

Independent claim 13 is directed to a node (*See e.g.*, FIG. 20, #140A) for use in a multi-node system (*See e.g.*, FIG. 20, #100). The node includes a plurality of devices including a system memory (*See e.g.*, FIG. 20, #144A), an active device (*See e.g.*, FIG. 20, #142A, 146A), and an interface (*See e.g.*, FIG. 20, #148A) configured to send and receive coherency messages on an inter-node network (*See e.g.*, FIG. 20, #154; specification page 63, lines 7-15) coupling nodes in the multi-node system. The node also includes an address network (*See e.g.*, FIG. 20, #150A) configured to convey address packets between the plurality of devices, and a data network (*See e.g.*, FIG. 20, #152A) that is separate from the address network and configured to convey data packets between the plurality of devices (*See e.g.*, specification page 11, lines 4-20). In response to receiving a coherency message on the inter-node network requesting an access right to

a coherency unit, the interface is configured to send a first type of address packet on the address network if the global access state of the coherency unit within the node is the modified state, and a second type of address packet if the global access state of the coherency unit in the node is not the modified state (*See e.g.*, specification page 92, line 24 - page 93, line 5). The system memory is configured to respond to receipt of the second type of address packet by sending a data packet to the coherency unit on the data network, regardless of whether the system memory currently has an ownership responsibility for the coherency unit (*See e.g.*, specification page 101, lines 11-15).

Independent claim 24 is directed to a method of operating a multi-node computer system (*See e.g.*, FIG. 20, #100) that includes a node (*See e.g.*, FIG. 20, #140A) and an additional node (*See e.g.*, FIG. 20, #140B) coupled by an inter-node network (*See e.g.*, FIG. 20, #154). The method includes an interface (*See e.g.*, FIG. 20, #148A) in the node receiving a coherency message requesting an access right to a coherency unit via the inter-node network from an additional interface in the additional node (*See e.g.*, specification page 63, lines 7-15). The method also includes the interface sending an address packet on an address network in the node in response to the receiving, wherein the address packet is a first type of address packet if the global access state of the coherency unit in the node is a modified state and a second type of address packet if the global access state is not the modified state (*See e.g.*, specification page 92, line 24 - page 93, line 5). The method further includes, in response to the sending, a system memory in the node (*See e.g.*, FIG. 20, #144A) providing, via a data network (*See e.g.*, FIG. 20, #152A) that is separate from the address network, the interface with data corresponding to the coherency unit regardless of whether the system memory has an ownership responsibility for the coherency unit if the address packet is the second type of address packet (*See e.g.*, specification page 101, lines 11-15).

VI. GROUND OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1-32 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Liencres et al. (U.S. Patent No. 5,434,993) in view of Chandrasekaran et al. (U.S. Patent No. 6,970,872) and in view of Roy (6,065,092).

VII. ARGUMENT

First Ground of Rejection:

Claims 1-32 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Liencres et al. (U.S. Patent No. 5,434,993) in view of Chandrasekaran et al. (U.S. Patent No. 6,970,872) and in view of Roy (6,065,092). Appellant traverses this rejection for at least the following reasons.

Independent claims (by number):

Appellant respectfully submits that each of claims 1, 13, and 24 recites a combination of features not taught or suggested in Liencres, Chandrasekaran or Roy. For example, claim 1 recites a combination of features including:

a node including an active device, a system memory, and an interface interconnected by an address network and a data network that is separate from the address network ...

wherein in response to the coherency message, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency unit in the node is a modified state and to send a second type of address packet on the address network if the global access state is not the modified state;

wherein in response to the second type of packet, the system memory is configured to send a data packet corresponding to the coherency unit on the data network, regardless of whether the memory has an ownership responsibility for the coherency unit.

Appellant's arguments are twofold: First, that Liencres does not teach or disclose "a node including an active device, a system memory, and an interface interconnected by an address network and a data network that is separate from the address network" as recited in Appellant's claim 1. Second, that Chandrasekaran does not teach or disclose "wherein in response to the coherency message, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency

unit in the node is a modified state and to send a second type of address packet on the address network if the global access state is not the modified state;” or “wherein **in response to the second type of packet**, the system memory is configured to send a data packet corresponding to the coherency unit on the data network, regardless of whether the memory has an ownership responsibility for the coherency unit” as recited in Appellant’s claim 1.

The Examiner asserts in the Advisory action dated March 18, 2008, and in the final Office action dated December 26, 2007 Liencres teaches Appellant’s claimed address network and a separate data network that couples the active devices, system memory, and the interface. More particularly, the Examiner asserts that bus 33 of Liencres is Appellant’s address network. Further, in the Office action dated July 10, 2007 the Examiner asserts “the memory and the active device are part of element 32 in Liencres, which is connected to element 33.” Appellant respectfully submits this is largely irrelevant.

As shown in FIG. 3a and 3b Liencres clearly shows that the memory alluded to by the Examiner is a cache memory 37, and not a system memory. In Appellant’s claims, specification and drawings it is clear the claimed memory is a system memory and not a cache memory. Appellant submits Liencres clearly shows a system (main) memory coupled by a memory bus 25 to each node. Notwithstanding, the Examiner still insists that the cache memory 37 is analogous to Appellant’s claimed memory. However, even if, *arguendo*, one were to suppose that the cache memory 37 were analogous to the Appellant’s claimed system memory, the topology of Liencres is different. Specifically, the cache memory 37 is not coupled to the bus 33 but to the cache controller 35 and to processor 21 by a separate bus. As such, the bus cache controller 31, the processor 21 and the cache memory 37 are NOT all coupled together by bus 33 as suggested by the Examiner. Thus, Appellant fails to see how the system of Liencres teaches the structure recited in Appellant’s claims. Appellant submits it does not.

In addition, the Examiner acknowledges Liencres does not teach “wherein in

response to the coherency message, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency unit in the node is a modified state and to send a second type of address packet on the address network if the global access state is not the modified state;" and "wherein in response to the second type of packet, the system memory is configured to send a data packet corresponding to the coherency unit on the data network, regardless of whether the system memory has an ownership responsibility for the coherency unit," as recited in Appellant's claims 1, 13, and 24. However, the Examiner asserts Chandrasekaran teaches the limitations at figure 1 and at col. 2, lines 54-57, col. 2, lines 60-62, and at col. 6, lines 25-36. Appellant respectfully disagrees.

More particularly, the Examiner asserts "Chandrasekaran teaches a multi-node network (figure 1) that employs several techniques to reduce latency. One of the methods is called "optimistic read" (col. 2, lines 54-57) where the system sends the read data regardless of whether or not the data is valid (i.e., modified) (col. 2, lines 60-62). If a request is made its validity is determined. A message is sent granting or denying access to the resource based on its validity. One of the methods of determining validity is "write-time" validity checking (col. 6, lines 25-36). When another node writes out data, it sends out a report stating the latest write time for that data. The read data is invalid once its timestamp comes before the latest write time. The node, now having an invalid read data, will have to request the updated data from the additional node. It would have been obvious... to employ optimistic reading of data using write time validity checking so that reads could be employed when another node has exclusive access but hasn't yet written the data." Appellant respectfully disagrees with the Examiner's application of the Chandrasekaran art to Appellant's claims.

More particularly, Chandrasekaran is directed to optimistic reads and write time validity checking. Taking Chandrasekaran in context, Chandrasekaran actually discloses

"This technique of starting the retrieval of the resource before receiving a response, such as a lock, to a request for permission to access the resource is referred to herein as an "optimistic read." The techniques described herein not only perform an optimistic read but also determine whether the

results of the optimistic read are valid, in the sense of providing the correct version of the resource retrieved. If the optimistic read is not valid, then the resource retrieved from the optimistic read is not used. In one embodiment of this aspect, if the version retrieved by the optimistic read is not valid, then another operation is initiated to retrieve the resource, but only after permission is received to access the resource. When the optimistic read results are valid sufficiently often, latency is reduced in retrieving resources.” (See col. 2 lines 54-67) (Emphasis added)

“In an embodiment using the first type of validity checking, the time that the optimistic read is started is compared to the latest time that the data block was written by any of the other nodes. If the read was started after the last write, the read is valid. **This can be determined even before the read is finished,** but involves the writing node publishing its write time to the other nodes. A node can publish its write time in any way, such as by broadcasting the write time to the other nodes, by storing the write time and responding to requests from other nodes, or by sending the write time to a lock manager. This type of validity checking is called “write-time” validity checking herein.” (See col. 6, lines 25-36) (Emphasis added)

From the foregoing, Appellant submits Chandrasekaran is disclosing an optimistic read operation in which the data is retrieved and then checked for validity. Appellant further submits the read data is apparently sent without regard to validity, and not whether the memory has ownership responsibility. This passage nor any other discloses “the system memory is configured to send a data packet corresponding to the coherency unit on the data network, regardless of whether the system memory has an ownership responsibility for the coherency unit,” as recited in claim 1.

Furthermore, there is no teaching of the interface sending one kind of packet if the coherency unit in the node is in a modified state and a second type of packet if it is not in a modified state. The memory then responds to the second type of packet as described above. But it is “regardless of whether the memory has an ownership responsibility for the coherency unit” and not whether or not the data is valid. These are distinctly different.

Thus, Appellant submits neither Liencres nor Chandrasekaran, taken either singly or in combination, teaches or suggests the combination of features recited in Appellant's claim 1.

Appellant's claims 13 and 24 recite features that are similar to the features recited in claim 1. Thus, for at least the above stated reasons, Appellant submits that the rejection of claims 1, 13, and 24 is in error and requests reversal of the rejection. The rejection of claims 2-12 (dependent from claim 1), claims 14-23 (dependent from claim 13), and claims 25-32 (dependent from claim 24) are similarly in error for at least the above stated reasons, and reversal of the rejection is requested. Each of claims 2-12, 14-23, and 25-32 recite additional combinations of features not taught or suggested in the cited art.

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-32 is erroneous, and reversal of his decision is respectfully requested.

The Commissioner is authorized to charge any fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5181-94901/SJC.

Respectfully submitted,

/ Stephen J. Curran /

Stephen J. Curran

Reg. No. 50,664

AGENT FOR APPELLANT

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: May 27, 2008

VIII. APPENDIX

The claims on appeal are as follows.

1. A system, comprising:

a node including an active device, a system memory, and an interface

interconnected by an address network and a data network that is separate from the address network;

an additional node coupled to send a coherency message to the interface in the node via an inter-node network, wherein the coherency message requests an access right to a coherency unit;

wherein in response to the coherency message, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency unit in the node is a modified state and to send a second type of address packet on the address network if the global access state is not the modified state;

wherein in response to the second type of packet, the system memory is configured to send a data packet corresponding to the coherency unit on the data network, regardless of whether the system memory has an ownership responsibility for the coherency unit.

2. The system of claim 1, wherein the coherency message requests a read access right to the coherency unit, wherein the first type of address packet is a proxy read-to-

share-modified packet and wherein the second type of address packet is a proxy memory read packet.

3. The system of claim 2, wherein if the active device has the ownership responsibility for the coherency unit, the active device is configured to send a data packet corresponding to the coherency unit to the interface via the data network in response to receipt of the proxy read-to-share-modified packet.

4. The system of claim 3, wherein if the active device has the ownership responsibility for the coherency unit, the active device is configured to lose its ownership responsibility for the coherency unit upon receipt of the proxy read-to-share-modified packet.

5. The system of claim 3, wherein if the active device has the ownership responsibility for the coherency unit, the active device is configured to transition an access right to the coherency unit upon sending the data packet on the data network.

6. The system of claim 1, wherein the address network is configured to convey the first and second types of address packet from the interface to a directory in point-to-point mode.

7. The system of claim 2, wherein the address network is configured to convey the first and second types of address packet from the interface to a plurality of devices

included in the node in broadcast mode, wherein the plurality of devices include the system memory and the active device.

8. The system of claim 1, wherein the data packet sent by the system memory includes an indication of the global access state of the coherency unit in the node.

9. The system of claim 1, wherein the coherency message requests a shared access right to the coherency unit.

10. The system of claim 1, wherein the additional node is configured to send the coherency message in response to an additional active device included within the additional node sending an address packet on an additional address network included within the additional node, wherein the address packet requests write access to the coherency unit, wherein the coherency unit is in a shared global access state in the additional node, and wherein the node is a home node of the coherency unit.

11. The system of claim 10, wherein if the coherency unit is in the shared global access state in any of the plurality of nodes other than the home node, the coherency unit is in the shared global access state in the home node and no active device and no memory subsystem included in any of the plurality of nodes has the ownership responsibility for the coherency unit.

12. The system of claim 11, wherein the interface is configured to send a copy of the coherency unit included in the data packet to the additional node.
13. A node for use in a multi-node system, the node comprising:
- a plurality of devices including a system memory, an active device, and an interface configured to send and receive coherency messages on an inter-node network coupling nodes in the multi-node system;
 - an address network configured to convey address packets between the plurality of devices;
 - a data network that is separate from the address network and configured to convey data packets between the plurality of devices;
- wherein in response to receiving a coherency message on the inter-node network requesting an access right to a coherency unit, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency unit in the node is a modified state and to send a second type of address packet on the address network if the global access state of the coherency unit in the node is not the modified state;
- wherein the system memory is configured to respond to receipt of the second type of address packet by sending a data packet corresponding to the coherency unit on the data network, regardless of whether the system memory currently has an ownership responsibility for the coherency unit.

14. The node of claim 13, wherein the coherency message requests a read access right to the coherency unit, wherein the first type of address packet is a proxy read-to-share-modified packet and wherein the second type of address packet is a proxy memory read packet.

15. The node of claim 14, wherein if the active device is the owner of the coherency unit, the active device is configured to send data corresponding to the coherency unit to the interface in response to receipt of the proxy read-to-share-modified packet.

16. The node of claim 14, wherein if the active device is the owner of the coherency unit, the active device is configured to lose its ownership responsibility for the coherency unit upon receipt of the proxy read-to-share-modified packet.

17. The node of claim 13, wherein the interface includes a global access state cache indicating global access states of a plurality of recently accessed coherency units in the node.

18. The node of claim 17, wherein the interface is configured to check the global access state cache for the global access state of the coherency unit in the node, wherein if the global access state of the coherency unit is not included in the global access state cache, the interface is configured to request an indication of the global access state of the coherency unit from the system memory.

19. The node of claim 18, wherein the interface is configured to request the global access state of the coherency unit in the node from the memory by sending the second type of address packet to the system memory.

20. The node of claim 13, wherein the data packet sent by the system memory includes a copy of the coherency unit.

21. The node of claim 13, wherein the data packet sent by the system memory includes an indication of the global access state of the coherency unit in the node.

22. The node of claim 13, wherein the address network is configured to convey the first and second types of address packet from the interface to a directory in point-to-point mode.

23. The node of claim 13, wherein the address network is configured to convey the first and second types of address packet from the interface to the plurality of devices in broadcast mode.

24. A method of operating a multi-node computer system, wherein the multi-node computer system includes a node and an additional node coupled by an inter-node network, the method comprising:

an interface in the node receiving a coherency message requesting an access right to a coherency unit via the inter-node network from an additional interface in the additional node;

the interface sending an address packet on an address network in the node in response to said receiving, wherein the address packet is a first type of address packet if the global access state of the coherency unit in the node is a modified state and a second type of address packet if the global access state is not the modified state;

in response to said sending, a system memory in the node providing, via a data network that is separate from the address network, the interface with data corresponding to the coherency unit regardless of whether the system memory has an ownership responsibility for the coherency unit if the address packet is the second type of address packet.

25. The method of claim 24, wherein the coherency message requests a read access right to the coherency unit, wherein the first type of address packet is a proxy read-to-share-modified packet and wherein the second type of address packet is a proxy memory read packet.

26. The method of claim 25, further comprising an active device included in the node sending data corresponding to the coherency unit to the interface in response to receipt of the proxy read-to-share-modified packet if the active device has the ownership responsibility for the coherency unit.

27. The method of claim 25, further comprising the active device losing the ownership responsibility for the coherency unit upon receipt of the proxy read-to-share-modified packet if the active device has the ownership responsibility for the coherency unit.

28. The method of claim 27, further comprising the address network conveying the first and second types of address packet from the interface to a directory in point-to-point mode.

29. The method of claim 24, further comprising the address network conveying the first and second types of address packet in broadcast mode.

30. The method of claim 24, further comprising the additional node sending the coherency message in response to an additional active device included within the additional node sending an address packet on an additional address network included within the additional node, wherein the address packet requests write access to the coherency unit, wherein the coherency unit is in a shared global access state in the additional node, and wherein the node is a home node of the coherency unit.

31. The method of claim 30, wherein if the coherency unit is in the shared global access state in any of the plurality of nodes other than the home node, the coherency unit is in the shared global access state in the home node and no active device and no memory

subsystem included in any of the plurality of nodes has the ownership responsibility for the coherency unit.

32. The method of claim 31, further comprising the interface sending a copy of the coherency unit included in the data packet to the additional node.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 C.F.R. §§ 1.130, 1.131, or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.